

UNCLASSIFIED

Defense Technical Information Center
Compilation Part Notice

ADP011126

TITLE: Practical Aspects of Implementing H-Infinity Controllers on a
FBW Research Helicopter

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Active Control Technology for Enhanced Performance Operational
Capabilities of Military Aircraft, Land Vehicles and Sea Vehicles
[Technologies des systemes a commandes actives pour l'amelioration des
performances operationnelles des aeronefs militaires, des vehicules
terrestres et des vehicules maritimes]

To order the complete compilation report, use: ADA395700

The component part is provided here to allow users access to individually authored sections
of proceedings, annals, symposia, etc. However, the component should be considered within
the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:
ADP011101 thru ADP011178

UNCLASSIFIED

Practical Aspects of Implementing H-Infinity Controllers on a FBW Research Helicopter

D.J. WALKER

M.C. TURNER

Department of Engineering
University of Leicester
Leicester LE1 7RH, UK

A.W. GUBBELS

Flight Research Laboratory
National Research Council of Canada, Bldg U-61
Ottawa, Ontario K1A 0R6

ABSTRACT

This paper presents a summary of the design and testing of two H-infinity controllers, recently flight-tested on the NRC's Bell 205 experimental fly-by-wire helicopter. Lessons learned from the implementation and testing are described. Both designs were based on low-order mathematical models and H-infinity optimization. The first controller successfully engaged first time, and is believed to be the first H-infinity controller flight-tested on a rotorcraft. It was subsequently evaluated at hover and low/moderate speed by a test-pilot, and found to achieve level 2 Cooper Harper Handling Qualities on a number of tasks. The controller was re-designed using a different mathematical model and a different H-infinity cost-function. The result was a significant reduction in cross-couplings, better (though still Level 2) handling qualities ratings of 4-5, Level 1 pitch and roll bandwidths. This paper presents an analysis of data from these flights. The flight testing provided a number of important practical lessons that could be useful to anyone attempting to implement and test modern controllers in flight. The gap between robustness of the design method and accuracy of the flight mechanic model is one of the most critical issues in high bandwidth control. Improved aircraft models translate directly into better controller performance. Validation of the aircraft model against open loop helicopter flight test data has shown that both the models used were deficient in a variety of ways. Software implementation should be kept as simple as possible; a discussion of the methods used for this project is given. The use of an on-board aircraft model greatly assisted in trouble-shooting the code for errors before flying. Use of automated code generation greatly reduces transfer errors from the Matlab design environment. To assess new control laws fully, an experienced test pilot is essential.

INTRODUCTION

Helicopter flight dynamics are governed by many complex and still quite poorly understood phenomena. This makes accurate modelling hard, and designing new and better controllers challenging. There is, furthermore, a need for better helicopter control systems: for example, to meet new and demanding specifications like ADS-33 (1989) that will be applied in future procurements, civil as well as military.

Interest in H. control of helicopters dates back to the work of Postlethwaite and his co-workers Tombs and Yue in the mid-late 1980's: see Tombs (1987); Yue, Postlethwaite and Padfield (1989); Yue and Postlethwaite (1990). The Westland Lynx was also the focus of a more recent study: see Walker & Postlethwaite (1990), (1996). This all suggested that H. design methods offered considerable promise in helicopter control, but the limitations of ground-based simulation were recognized, and it was concluded by Walker & Postlethwaite (1996) that appropriately validated mathematical models containing higher order dynamics would be important if similar (Level 1) results were to be replicated in flight.

Recent collaboration between the University of Leicester and the National Research Council (NRC) of Canada has resulted in the opportunity to test H. controllers in flight using NRC's modified Bell 205 helicopter. This paper presents analysis of data from that work, and a comparison between two designs, which were tested in August 1997 and February 1999. More details of the respective flight tests are given by Postlethwaite *et al* (1998) and Walker *et al* (1999) respectively.




Figure 1: Bell 205 Airborne Simulator

BELL 205 AIRBORNE SIMULATOR

The Bell 205 is a multi-role utility and transport helicopter. NRC's Bell 205 airborne simulator (Figure 1) is an extensively modified version of the Bell 205A-1: see Sattler, (1984). Amongst the modifications, the standard Bell 205 stabilizer bar has been removed to enhance the control response of the teetering rotor. The aircraft serves as a fly-by-wire variable stability platform for in-flight simulation of other aircraft, and investigation of control systems and new cockpit technologies. It is configured to have a Safety Pilot (SP) flying from the left-hand seat and an Evaluation Pilot (EP) flying from the right-hand seat. The original actuators have been replaced with specially built dual-mode electro-hydraulic actuators that can be either electrically or mechanically controlled. During normal flight, these actuators are mechanically controlled and behave just as the original actuators. During FBW flight, they are electrically controlled by the EP, but can be mechanically overridden if the SP supplies a force that exceeds a pre-set control breakout force. During the H. controller tests, both EP and SP were test-pilots. The EP's role was to evaluate the experimental controller. The role of the SP is to monitor the aircraft actuator control activity fed back to the SP cyclic and commanded by the experimental controller. If this control activity is oscillatory, divergent or otherwise overly active, he will disengage the FBW system and take control of the aircraft. The aircraft has equipment for measuring and recording many variables, including: 3-axis attitudes, angular rates, accelerations, velocities, static and dynamic pressure, air temperature, angle-of-attack, side-slip, pilot control inputs and actuator positions.

AIRCRAFT MODELS

Two different mathematical models were used to design the controllers. These are now described.

NASA Model The linear model used to design the first controller (henceforth referred to as Controller I) was a 10 knot, six-state stability and control derivative model from a NASA contractor report by Heffley *et al* (1979). The six states are the three angular velocity components of the fuselage (p, q, r) and three translational velocity components of its mass centre (u_B, v_B, w_B). This model was augmented with pitch and roll attitudes (θ and ϕ) to enable an Attitude Command/Attitude Hold (ACAH) response to be designed. The dynamic response of a teetering rotor is sometimes represented by a transport delay, so the nominal model was cascaded with first-order Padé approximants in each of the three control channels. This doubled as a simple actuator model. The Padé approximation time constants were chosen to reflect the effective combined actuator and rotor time constants: i.e. about 0.156 sec in pitch and roll and 0.187 sec in yaw actuators. The first controller design was based on the resulting model.

Using a low-order rigid-body model for control law design has potential drawbacks. The NASA model captures the salient rigid body modes reasonably well, but the omission of rotor dynamics limits the achievable bandwidth. Furthermore, the stability derivatives given by Heffley *et al* relate to a standard model 205 with a stabilizer bar. The bar has been removed from the NRC Bell 205 on which the flight tests were conducted.

DERA Helisim Model (Padfield, 1981.) The Helisim generic helicopter non-linear model has been used for over ten years in its various Lynx configurations as the basis for the design and simulation of H. and other novel controllers. After Controller I had been tested, DERA re-configured the Helisim model to represent the NRC Bell 205; see Strange & Howitt, (1997). DERA also undertook a validation exercise on the Helisim 205 model, using flight test data from NRC's Airborne Simulator. Comparison was also made with the NASA linear model (case 126) from Heffley *et al.* (1979). Although the control law design work was conducted at hover/low-speed, the model validation was performed at 60 knots because at the time of model development no high quality open-loop flight test data were available for the hover condition. The validation exercise indicated:

- Using a quasi-steady approximation to the rotor dynamics, the model captures the basic rigid body behaviour reasonable well, at least on-axis.
- Fidelity of the model can be further enhanced by incorporating coning and flapping dynamics, an inflow correction factor, and tail fin blockage and tail rotor blade root cut-out effects.
- Significant model uncertainty still exists and the model can only be considered to be of low/medium fidelity.
- The Helisim 205 model generally gives a better correlation than does the NASA model with flight data from NRC's Bell 205.

A nineteen-state 20-knot linearization from the Helisim model was the basis for the second design (Controller II) discussed here.

CONTROLLER DESIGN

Both controllers were designed to give an Attitude Command/Attitude Hold response type in pitch and roll, and a rate command in yaw. Both used five measurements: pitch

and roll attitudes and rates and yaw rate, and had full authority over three control inputs: longitudinal and lateral cyclic and tail-rotor collective. Main rotor collective was left open loop, partly because the available instrumentation provided no suitable heave-axis velocity measurement with which to close that loop.

CONTROLLER I: MULTIVARIABLE LOOP SHAPING

The first design followed in precisely the same manner as the Lynx design of Walker and Postlethwaite (1996). It was based on the two degree-of-freedom H_∞ optimization proposed by Hoyle *et al* (1991). The main steps in the design process are: (i) selection of a low-order step response model (SRM) that encapsulates basic handling requirements; (ii) augmentation of the aircraft $G(s)$ at input and output with filters W_1 and W_2 ; (iii) synthesis of a stabilizing controller $K(s)$ minimizing the H_∞ norm of the transfer function from $\{v, w\}$ to $\{u, y, z\}$ (see Figure 2); (iv) incorporation of filters into $K(s)$. Note that while only three outputs were actually controlled, measurements of a further two (the rates q and p) were also fed back into the controller, to enhance stability. H_∞ optimization produces a controller that forces the closed loop to approximate the SRM, by reducing the H_∞ norm of the difference between the two.




Figure 2: Two DoF Structure: Controller I

The design is based around a normalized coprime factor description of the augmented nominal aircraft model. $[N \ M]$ denotes a left coprime factorization of the nominal augmented aircraft transfer function matrix $G(s)$. This means that $G = M^{-1}N$, in which N and M are stable and there is no cancellation of any unstable dynamics between M^{-1} and N . The factorization is said to be *normalized* if, in addition, $[N \ M]$ is all-pass. Further relevant information is given by McFarlane and Glover (1990).

Step Response Model A second order system with no zeros and with damping ratio ζ and undamped natural frequency ω_0 was used in each of the three controlled axes. Damping ratios and natural frequencies are given in Table 1.

	Roll	Pitch	Yaw
Damping ratio (ζ)	0.9	0.9	0.7
Nat. Freq. (ω_0) (rad/s)	3.6	1.2	11.0

Table 1: Step Response Model Parameters

Loop Shaping The nominal aircraft model was pre- and post-multiplied by filters $W_1(s)$ and W_2 to produce a singular value loop-gain consistent broadly with frequency-domain performance and robustness requirements: high d.c. gain for steady-state disturbance rejection and tracking; low gain at high frequencies for robustness, noise attenuation etc.; slope

not less than minus 40 dB/decade at gain cross-over for stability. The filters chosen were:

$$W_1 = \text{diag}\left(\frac{s+2}{s}, \frac{s+2}{s}, \frac{s+2}{s}\right)$$

$$W_2 = \text{diag}(1, 1, 1, 0.5, 0.5)$$

To set the 0 dB crossover and to help decouple the aircraft in the mid frequency range, alignment was performed at $\omega = 4$ rad/s.

Design Optimization The controller K achieving the desired performance in the face of the assumed model error is obtained by minimizing the H_∞ norm of the closed loop transfer function T from $\{v, w\}$ to $\{u, y, z\}$ subject to internal stability: see Figure 2. This has the effect of simultaneously reducing the energy in u, y , and z due to commands (v) and model error (w). The parameter ρ allows model-following to be balanced against robust stabilization; setting $\rho=0$ one reverts to the pure robust stabilization problem of McFarlane and Glover (1990). $\rho = 1.9$ gave a satisfactory compromise.

Stabilizing controllers minimizing the appropriate H_∞ norm can be found using standard algorithms for H_∞ optimization. Alternatively, as was the case here, an observer-based controller satisfying the above H_∞ optimization criterion can be found directly from formulae given by Walker (1996). This structure was exploited in the C-code implementation of the controller for real-time implementation. Euler numerical integration was used to integrate all the dynamic subsystems within the controller: i.e. the observer, ideal model and filter W_1 . Also included in the command path were a dead-band and a low-pass filter.

CONTROLLER II: DECOUPLED MIXED-SENSITIVITY (S/KS) OPTIMIZATION

This design was based on the DERA Helisim model. The aim was to investigate whether robustness could be improved by eliminating from the control-law design terms representing cross-couplings. The resulting Controller II consisted of two independent sub-controllers for longitudinal and lateral dynamics respectively, each based on an appropriate decoupled model.

Longitudinal and Lateral Models The Helisim 205 model was linearized about a twenty knot flight condition to yield a nineteen-state linearization: nine rigid body, six rotor dynamic, and four actuator states. Based on the premise that Helisim's rotor representation was reliable at predicting steady-state rotor forces but less reliable in terms of its transient predictions, the six model states corresponding to rotor flap and coning were residualized: i.e. assumed to reach their steady-state values instantaneously. The resulting model was partitioned into longitudinal ($\theta, q, u_B, w_B, \eta_{1s}$) and lateral ($\phi, p, r, v_B, w_B, \eta_{1c}, \eta_{0t}$) states for design of the respective controllers (η 's represent actuator states).

Figure 3 shows the structure used with Controller II, viz. separate longitudinal and lateral controllers. A full authority Attitude Command/Attitude Hold (ACAH) response type was again specified, in which pitch and roll attitudes and yaw rate (θ, ϕ, r) are demanded by the pilot. Pitch and roll rates (q and p) are again fed back in their respective loops. The longitudinal controller controls longitudinal cyclic (θ_{1s}) while the lateral controller controls lateral cyclic (θ_{1c}) and tail rotor collective (θ_{0t}). The main rotor collective was again left open loop. The longitudinal and lateral controllers were designed separately using continuous-time methods, then

converted to discrete-time equivalents and implemented in state-space form at a fixed sample-rate of 64 Hz on the aircraft's flight-control computer. Dead-band and low-pass filters were again employed in the command path in order to clean up the signals from the pilot inceptors.

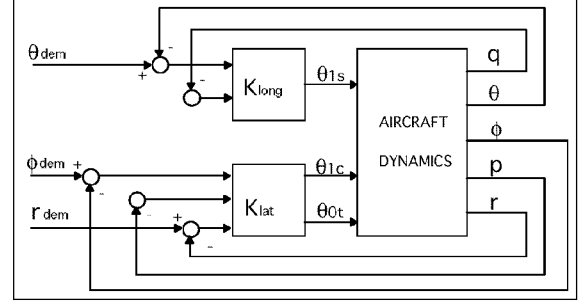


Figure 3: Decoupled Structure of Controller II

Design Optimization A weighted mixed-sensitivity, or S/KS, H_∞ optimization was used in both longitudinal and lateral designs. A stabilizing controller $K(s)$ is sought such that:

$$K(s) = \arg \min \left\| \begin{bmatrix} W_1(I + GK)^{-1} \\ W_2K(I + GK)^{-1} \end{bmatrix} \right\|_\infty$$

The above cost function leads to a stabilizing controller that simultaneously attempts to reduce the energy in the weighted tracking-error and weighted control signal due to commands or disturbances at the aircraft output. The principal design parameters were a sensitivity weight W_1 and a robustness weight W_2 . W_1 was a low-pass filter, used to shape the sensitivity function $(I + GK)^{-1}$; its 0 dB cross-over approximately defines the tracking bandwidth. W_2 is a high pass filter, used to shape $K(I + GK)^{-1}$, which in turn governs robustness to additive model error and control usage.

After some iteration, the filters chosen were:

$$\text{Pitch Attitude: } W_1 = \frac{0.5}{s + 0.01}$$

$$\text{Roll Attitude: } W_1 = \frac{0.3s + 0.15}{s + 0.01}$$

$$\text{Yaw Rate: } W_1 = \frac{0.3s + 0.15}{s + 0.01}$$

$$\text{Pitch and Roll Rates: } W_1 = \frac{s}{s + 0.01}$$

$$\text{Control Weight: longitudinal cyclic: } W_2 = \frac{2s + 0.002}{s + 5}$$

Control weight: lateral cyclic and tail rotor collective:

$$W_2 = \frac{2s + 0.002}{s + 4}$$

Solving two separate H_∞ optimizations using standard software algorithms led to longitudinal and lateral controllers of 8 and 12 states respectively. These were discretized using a zero-order hold for digital implementation.

MIXED ANGULAR RATES

Experience at NRC has been that feeding back measured rate signals p and q can lead to stability problems. This is believed to be due in part to the fact that the angular rate sensors detect structural modes as well as rigid body motion.

More important still, the teetering rotor interposes a characteristic dead time of between 120-180 mS in the control responses of the various axes: a dead-time that the models do not predict properly. Use of so-called mixed rate signals to alleviate this is described by Baillie *et al* (1994). The mixed rate signal is synthesized using a pair of complementary filters; the measured rate is fed to the low-pass filter, then summed with the output of the high-pass filter. The latter is driven by an open-loop predictor that consists simply a lag-free first-order model of the on-axis control response, driven in turn by the actuator command. As implemented on the NRC Bell 205, the mixed rate feedback is essentially an open-loop predictor-derived signal at frequencies above approximately 11.0 rad/s. Controllers I and II were both driven by mixed rate signals.

RESULTS

Controller I Figure 4 shows the primary response to a doublet demand on pitch attitude. The response lags by approximately 2 seconds. Pitch response was designed to be slower than roll (not shown), to allow for the aircraft's greater inertia about its pitch axis. The roll response (not shown) was damped oscillatory. Although the aircraft was quite flyable, undesirable cross-couplings were present, and the bandwidths achieved were quite low: considerably lower than predicted: see Table 2. (Bandwidths predicted using small-signal analysis are shown in parentheses).




Figure 4: Controller I Pitch Attitude Response

Frequency sweeps were also conducted during the flight. Spectral analysis was used to extract frequency responses from the sweep data. Bode plots for the pitch axis are shown in Figure 5. Coherence of greater than about 0.7 is generally taken to indicate reasonable frequency response identification. Gain and phase information can be used to calculate the handling qualities bandwidth and phase delay parameters defined in ADS-33, and discussed by Padfield (1996). For an attitude command/attitude hold response type, the handling qualities bandwidth is defined as the frequency at which the phase of the closed loop system equals -135° . Phase delay is a measure of the rate-of-change of phase with frequency beyond the crossover frequency. It is defined as the ratio of the additional phase lag (in radians) beyond $-\pi$ rad at twice the bandwidth frequency to twice the bandwidth frequency (in rad/s).




Figure 5: Controller I Pitch Axis Frequency Response

Bandwidth and phase delay for Controller I are represented by the '*' in Figures 8 and 9. In terms of short-term frequency response criteria, Controller I is Level 3 for combat/target tracking and Level 2 for all other mission task elements. The Cooper-Harper Levels and Pilot Ratings are explained in Figure 10.

Controller II The doublet response on pitch axis for controller II is shown in Figure 6. The EP reported that inter-axis coupling was not an issue at all with this controller. Frequency response (pitch axis) is shown in Figure 7. The bandwidth and phase delay parameters are represented by the '+' in Figures 8 and 9.




Figure 6: Controller II Pitch Attitude Response

This controller resulted in a Level 1 system in terms of its pitch and roll short term frequency response at hover/low speed. Overall it was rated Level 2 because of deficiencies, the principle being a yaw response that was too slow and unpredictable; pitch response was also deemed to be a bit too sluggish. Roll axis response was deemed to be about right.

Both controllers were subjected to rigorous testing in a set of ADS-33 manoeuvres. The procedures are described in more detail by Postlethwaite *et al* (1998) and Walker *et al* (1999). The test pilot ratings are summarized in Table 3.




Figure 7: Controller II Pitch Axis Frequency Response




Figure 8: Bandwidth and Phase-Delay (Pitch)

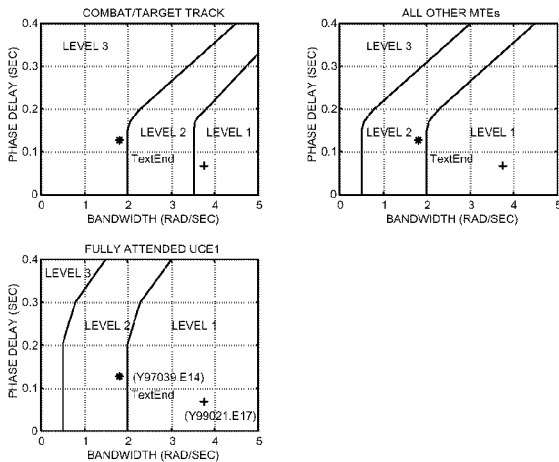


Figure 9: Bandwidth and Phase-Delay (Roll)

	Controller I		Controller II	
	Pitch	Roll	Pitch	Roll
B/W (rad/s)	1.32 (1.83)	1.81 (3.18)	2.64 (1.83)	3.75 (2.26)
Phase delay (s)	0.20	0.13	0.08	0.07

Table 2: Achieved (Predicted) bandwidths

Task	Controller I	Controller II
quick-hop	5	4
side-step	5	4
turn-to-target	4	5
precision hover	4	4
pirouette	5	4

Table 3: Handling Qualities Ratings from Flight-Tests

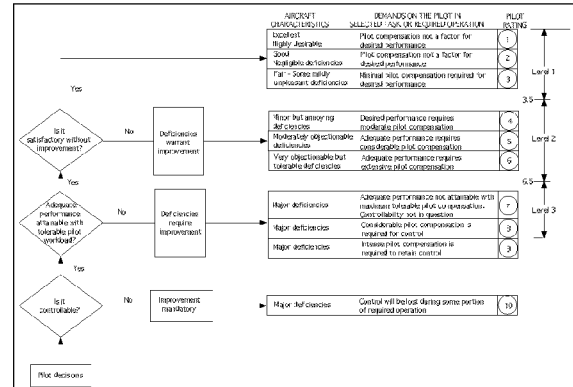


Figure 10: Cooper-Harper Rating Scheme

HQR 4: Minor but annoying deficiencies; Desired performance requires moderate pilot compensation;

HQR 5: Moderately objectionable deficiencies; Adequate performance requires considerable compensations;

LESSONS LEARNED

During the course of the H-infinity testing, many valuable lessons were learned that will help make future tests more efficient. These lessons would also be useful to any one attempting to implement and flight-test a modern controller on an aircraft.

Model Fidelity Two different aircraft models were used for controller designs in this experiment, as described above. Readily apparent deficiencies in the NASA model lead to the development of the DERA Helisim model. Validation of the response of both of these aircraft models against open loop helicopter flight test data has shown that both the models were deficient in a variety of ways. However, the DERA Helisim model was of higher fidelity than the NASA model. The result was that controllers designed with the Helisim model were more than twice as likely to function more-or-less as intended than were those designed with the NASA model. Another observation was that although the use of mixed rates was always required for controllers designed with the NASA model, not all of the controllers designed with the Helisim model required their use. One controller that was tested worked reasonably well, but only when a pre-filter was inserted into the yaw axis control input to smooth out the response. This type of remedy would not have been required had the aircraft model captured all the required dynamics adequately. It is for this reason that future control work is planned on the NRC Bell 412 Advanced Systems Research Aircraft, with this work being preceded by a model development program designed to create a full-envelope model of the aircraft that is of the highest possible fidelity.

Software Implementation The method of software implementation for the H-infinity program was not static from the first flight tests, and evolved a number of times. Since these controllers involved a large amount of matrix

manipulation and integration at each time cycle, processing time had to be carefully monitored to ensure that cycles were not overframing. Therefore, the code was kept as simple as possible, and was written in-line rather than in function calls to optimize for speed. Although initially the switch from ground-test mode to flight-test mode required a re-compile of the code, this was streamlined by adding a function switch to perform this change on-line. The switch was only active before a fly-by-wire engagement to prevent a change in mode while the evaluation pilot was flying the controller. For the first NASA-model based controllers, integrations were performed on-line in the code using a simple Euler method for faster computations. However, when the Helisim based controllers were implemented, the Euler method of integration was found to no longer be adequate, and was prone to instability. Several different methods of integration were attempted, but were either also prone to instability or required too much computational time, causing the flight control computer to skip cycles. The solution was found in using discretized controllers, which worked flawlessly and required much less computational time.

Automated Code Generation During the course of the first H-infinity experiment, it was quickly realized that the process of copying and implementing a large number of controllers, each consisting of 6 or more large matrices, from the Matlab environment to the Flight Control Computer was cumbersome. The solution was to write a Matlab routine that automatically generated C code from the output of the controller design process. It was then a simple matter to transfer the new controller file down to the helicopter and implement the new system. It is estimated that this automatic code generating routine required approximately 2 hours to write and saved over 8 man-hours of manual code manipulation per experiment. The other benefit of this system was that it was far less prone to coding errors than the manual method.

On-board Aircraft Model During the early stages of the work, when several controllers did not function as intended, it was unclear if this was as a result of implementation (e.g. coding) errors, or poor robustness of the design. Unlike with some standard controllers that will give meaningful output when tested open-loop, the H-infinity controllers, with a high integral component, cannot be tested open-loop on the helicopter prior to flight. A solution was devised to help eliminate implementation errors in the controller C-code; an aircraft linear model was implemented in the flight control computer and linked to the controller code. In this way the model was used to provide a closed-loop simulation of the aircraft/controller system in real-time using the same controller code. This greatly assisted in trouble-shooting the code for errors before flying, since this hardware-in-the-loop simulation proved that code was implemented correctly. If the controller did not work properly with the on-board model, then it was certain to fail in flight. (The converse is obviously not in general true, owing to the inevitable mismatch between the dynamics of the model and of the actual aircraft).

Experienced test pilot From past experience it was known that the use of operational evaluation pilots for handling qualities investigations is not desirable, and that the use of a qualified test pilot is essential to getting useful comments on the handling of the aircraft/controller system. Operational pilots are not generally capable of assigning appropriate Cooper-Harper ratings, because they have not been trained in the use of the scale: see Cooper & Harper (1969). Experience with the H-infinity flight tests have also shown that even trained, qualified test pilots often have difficulty defining

deficiencies with the control system if they have had no prior experience in performing handling qualities evaluations of control systems. During the trials, the NRC test pilot (who has several years experience working in the area of handling qualities of fly-by-wire control systems with a variety of response types) was able to provide excellent handling information for directing new controller designs.

RECOMMENDATIONS FOR FUTURE WORK

Future work will concentrate on refinement of the H-infinity designs and on application of different synthesis methods, principally QFT.

Future rotorcraft models at NRC will include rotor states. This is almost certainly necessary for designing high-bandwidth control systems. Future work will include rotor-state-feedback controllers designed using these models.

CONCLUSIONS

Two helicopter controller designs have been compared. The first design clearly demonstrated the potential of H_∞ optimization in the field of helicopter control. The controller led to a stable, flyable system. However, the first design caused us to ask whether, given current model fidelity, it would be better to decouple longitudinal and lateral controls, thereby reducing our reliance on a dubious part of the design model. Despite the yaw axis response of the second design being somewhat worse, the changes resulted overall in a number of significant improvements, particularly in reduced cross-coupling, so we would tentatively answer the above question in the affirmative.

The second design led to a stable and controllable system with desired performance, albeit with moderate work-load, in which coupling was not an issue. Predicted and achieved bandwidths showed fair agreement. Modifications to the controller design parameters had the desired effect on closed-loop responses. This leads us to believe that Level 1 handling qualities (satisfactory without improvement, desired performance requiring minimal pilot compensation) will be achievable with this type of controller.

Neither controller functioned without mixed rates, but we have demonstrated that it is possible to use a theoretically-derived flight mechanic model as the basis for multivariable controller design. This has important implications for the development of future generation flight control systems.

Practical lessons learned from this research include:

- Models of the highest fidelity are a requirement for the design of high bandwidth modern controllers.
- Software implementation issues are important, particularly for control of the flow of calculation and in deciding on how to perform integrations.
- Automatic code generation saves time and is less prone to errors than manual code manipulation.
- An on-board test facility to determine if code has been correctly implemented on the aircraft prior to flight is a great asset.
- The use of an experienced test pilot for handling qualities evaluations is essential.

ACKNOWLEDGEMENTS

The work was funded by the UK Engineering and Physical Sciences Research Council. The contributions of Alex Smerlas (now at Eurocopter Deutschland) and Michael Strange and Jeremy Howitt (DERA Bedford) to the design of Controller I and the model development, respectively, are gratefully acknowledged.

REFERENCES

- Anonymous. 1989. *Handling quality requirements for military rotorcraft ADS-33C*. US Army Aviation Systems Command.
- Baillie, S.W., Morgan, J.M. and Goheen, K.R., 1994. *Practical experiences in control system design using the NRC Bell 205 airborne simulator*. AGARD Flight Mechanics Panel Symposium on Active Control, Turin, Italy, May 1994.
- Cooper G. E., and Harper, R.P., *The use of pilot rating in the evaluation of aircraft handling qualities*, National Aeronautics and Space Administration, Technical Note TN D-5153, Washington D.C., April 1969
- Heffley, R.K., Jewell, W.F., Lehman, J.M., and Van Winkle, R.A. 1979 *A compilation and analysis of helicopter handling qualities data*. Contractor report 3144, NASA.
- Hoyle, D., Hyde, R. and Limebeer, D., 1991. *An H_∞ approach to two degree of freedom design* Proc IEEE 30th C.D.C. pp 1581-1585 Brighton U.K.
- McFarlane, D. and Glover, K., 1990, *Robust controller design using normalized coprime factor plant descriptions*. Springer Verlag Lecture notes in control and information science series.
- Padfield, G. D., 1981. *A theoretical model of helicopter flight mechanics for application to piloted simulation*. Technical Report RAE TR 81048, Royal Aircraft Establishment, Farnborough.
- Padfield, G.D. 1996. *Helicopter flight dynamics*. AIAA Education Series.
- Postlethwaite, I., Smerlas, A. Walker, D.J., Gubbels, A.W., Baillie, S.W., Strange, M.E., Howitt, J., Horton, R.I., 1998 *H_∞ control: From desk-top design to flight test with handling qualities evaluation*. Annual Forum Proceedings - American Helicopter Society, v 2, p 1313-1324
- Sattler, D.E., 1984. *The National Aeronautical Establishment airborne simulator facility*. NRC, NAE Misc 58, May 1984.
- Strange, M.E. and Howitt, J., 1997. *Configuration of the DERA HELISIM model to represent the flight dynamics of the NRC Bell 205 fly-by-wire research helicopter*. DERA Technical report AS/FMC/TR97459/1, March 1997.
- Tombs, M.S. 1987. *Robust control system design with application to high performance helicopters*. D.Phil Thesis, University of Oxford.
- Walker, D. J. and I. Postlethwaite. 1990. *Full authority active control system design for a high performance helicopter*. Proc. XVI European rotorcraft forum., III.3.2.1-III.3.2.14.
- Walker, D.J. & Turner, M.C, Smerlas, A.J., Strange, M. and Gubbels, A.W. 1999. *Robust control of the longitudinal & lateral dynamics of the Bell 205*. Proc. 1999 American Control Conference pp 2742-2746.
- Walker, D.J. and Postlethwaite, I., 1996. *Advanced helicopter flight control using two-degree-of-freedom H_∞ optimization*. Journal of Guid., Cont., and Dynamics, Vol. 19, No.2, pp.461-468.
- Walker, D.J., 1996. *On the structure of a two-degree-of-freedom H_∞ loop shaping controller*. International Journal of Control 63, No. 6, 1105-1127
- Yue, A. and I. Postlethwaite. 1990. *Improvement of helicopter handling qualities using H_∞ optimization*. IEE Proceedings Part D 137:115-129.
- Yue, A., Postlethwaite, I., Padfield, G. 1989. *H_∞ design and the improvement of helicopter handling qualities*. Vertica, v 13, n 2, p 119-132.

Paper #31

Q, by David Moorhouse: In simple terms, you designed to desired responses. You should be able to get them with any methodology. Can you say why you did not achieve them and what you plan to do differently with your H-infinity refinements?

A: (Daniel Walker's): The performance was indeed not robust. Achieving robust performance when the dynamics are complex and the model poor is difficult. Trying to identify reasonable bounds on the various 'deltas' is something that needs more work. The available models are better in some respects than in others. For instance, they don't predict pitch-roll coupling well, so it makes sense not to make the controller design overly reliant on that part of the model. So, how best to use existing models is something I will concentrate on. Getting better models (and that probably means higher-order ones, with rotor and other modes) is also vital.

We are still (as a community) in the relatively early stages of devising ways to employ the degrees-of-freedom that H-infinity, mu-synthesis provide.

Q, by P.M Lodge: Do you have a feel for why the yaw controller bandwidth reduced in moving from controller I to controller II?

A: (Daniel Walker's): Controller I performed well in yaw. That perhaps led us to a false sense of security, assuming yaw to be the easy one. With hindsight, we didn't analyze the yaw behaviour carefully enough beforehand. We were expecting problems with pitch and roll, not yaw. (Subsequent analysis did reveal the low yaw bandwidth, but how much heed we would have paid to this, I'm not sure, because predictions based on simulations did not always prove reliable, to say the least!) An additional factor was that we simply had far less experience with the optimization used to synthesize controller II.

The Bell 205 apparently has a slightly nasty, second-orderish yaw mode, possible related to structural flexing. I'm fairly certain that none of the models we have captures that effect. All-in-all, the yaw axis control problem isn't trivial, and given the quite poor models we have, there was probably an element of hit-and-miss.